arXiv:1511.04745v1 [stat.ME] 15 Nov 2015

Model Space Priors for Objective Sparse Bayesian Regression

Andrew J. Womack†

Indiana University, Bloomington, IN, USA.

Claudio Fuentes

Oregon State University, Corvallis, OR, USA.

Daniel Taylor-Rodriguez†

SAMSI and Duke, Durham, NC, USA.

Summary. This paper investigates the construction of model space priors from an alternative point of view to the usual indicators for inclusion of covariates in a given model. Assumptions about indicator variables often lead to Beta-Binomial priors on the model space, which do not appropriately penalize for model complexity when the parameters are fixed. This can be alleviated by changing the parameters of the prior to depend on the number of covariates, though justification for this is lacking from a first-principles point of view. We propose viewing the model space as a partially ordered set. When the number of covariates increases, an isometry argument leads to the Poisson distribution as the unique, natural limiting prior over model dimension. This limiting prior is derived using two constructions that view an individual model as though it is a "local" null hypothesis and compares its prior probability to the probability of the alternatives that nest it. We show that this prior induces a posterior that concentrates on a finite true model asymptotically. Additionally, we provide a deterministic algorithm that takes advantage of the nature of the prior and explores good models in polynomial time.

1. Introduction

This paper considers models of the form

$$\mathbf{y} = \mathbf{X}_0\boldsymbol{\beta}_0 + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon} \quad \text{where} \quad \boldsymbol{\epsilon} \sim N_n(\mathbf{0}, \mathbf{I}/\tau). \tag{1}$$

In (1), $\mathbf{X}_0$ corresponds to covariates that *must* be used to model the mean structure of $\mathbf{y} = (y_1, \ldots, y_n)$; it contains p_0 covariates, one of which must be $\mathbf{1}$ and represents the intercept term. The matrix $\mathbf{X}$ corresponds to p covariates whose importance is to be tested. The model M_A for $A \subseteq \{1, \ldots, p\}$ corresponds modeling the mean structure as $\mathbf{X}_0\boldsymbol{\beta}_{0,A} + \mathbf{X}_A\boldsymbol{\beta}_A$ where $\mathbf{X}_A$ is the

†Funded by the National Science Foundation DMS Grant 1105127.

submatrix of $\mathbf{X}$ containing the columns indexed by $a \in A$. The precision in model M_A is denoted by τ_A and M_A comes equipped with a prior distribution $\pi(\boldsymbol{\beta}_{0,A}, \boldsymbol{\beta}_A, \tau_A | M_A)$. The model space $\mathcal{M}$ consists of all models M_A for $A \subseteq \{1, \ldots, p\}$ and a model M_A is nested in M_B ($M_A \subset M_B$) whenever $A \subsetneq B$.

Estimation and model selection in this framework are among the oldest statistical problems considered. This model is experiencing a recent resurgence as a key statistical framework due to problems where p is large and increasing with n (especially the case when $p >> n$). The model selection problem is exponentially complicated in p (due to the 2^p models in the model space). We address the selection problem through the specification of a prior distribution on the model space $\pi_p(M_A)$, while "objective" priors are used for the model specific parameters. Many important papers address the selection problem from the Bayesian framework, recommending both priors on the model space and priors on the model specific parameters.

Priors on the model specific regression coefficients have been extensively discussed in the literature. Many of the priors are a modification of Zellner's g-prior (Zellner, 1986). Zellner and Siow (1980) used Cauchy priors, which were originally advocated for by Jeffreys, and Liang et al. (2008) developed alternative mixtures of g-priors. The intrinsic prior's (Berger and Pericchi, 1996) model selection properties have been extensively investigated (Casella et al., 2009; Moreno et al., 2010; Girón et al., 2010; Casella and Moreno, 2012). Womack et al. (2014) showed that the intrinsic prior is a scaled mixture of g-priors and compared the intrinsic prior to other scaled mixtures of g-priors. Laplace priors have also been considered (Hans, 2010; Castillo et al., 2015). One advantage of Laplace priors over scaled mixtures of g-priors is that the latter usually depend on the observed design matrix $\mathbf{X}$, while the former can be applied to situations where $(\mathbf{X}_0\ \mathbf{X})$ is not full rank (for instance, when $p > n$). An alternative to these "local" priors (priors that are positive and continuous at the origin) are non-local priors (Johnson and Rossell, 2010, 2012; Altomare et al., 2013), which converge to zero at the origin. In general, local priors provide a $n^{d/2}$ learning rate against d false positives while non-local priors can provide far faster learning rates (powers larger than $d/2$). When p is growing with n, non-local priors provide one route towards consistent model selection. The other route is to create additional penalization through the prior on the model space.

As many authors have noted, the prior on the model space becomes very important when p is growing with n. Ley and Steel (2009) and Scott and Berger (2010) advocate for Beta-Binomial(a, b) priors on indicator variables for covariate inclusion. These authors advocate for $a = b = 1$, which provides a Bayesian version of a multiplicity correction. Unfortunately, as Johnson and Rossell (2012) notes, this multiplicity correction is not enough to provide posterior concentration on a

finite true model when local priors are used, providing an argument in favor of non-local priors. The issue is related to prior concentration in the model space; Beta-Binomial priors concentrate all of their mass on models with a fraction of p covariates being utilized, which becomes problematic as p increases. In order to address this issue, Wilson et al. (2010) use Beta-Binomial priors where $a = 1$ and $b = \lambda p$, which provides for a finite prior expected number of included covariates as p increases. Finally, Castillo et al. (2015) obtain beautiful selection and reconstruction results by requiring sightly faster than exponential decrease in prior probability of the number of covariates in the model. In particular, this can be achieved by taking $b = p^u$ for some $u > 1$. However, outside of constructions using infinite exchangeability, there is little *a priori* justification for priors on the model space.

We propose new criteria for developing priors on the model space. In Section 2, we show that the Poisson distribution for the model dimension provides the unique family of distributions satisfying a natural assumption on the model space. We view each model as a "local" null hypothesis and require that a model's prior probability is proportional to the prior probability of the set of models that nest it (these constitute the set of "local" alternatives"). The derived prior is compared to the priors previously developed (Scott and Berger, 2010; Wilson et al., 2010; Castillo et al., 2015) in terms of their penalization of model dimension. In Section 3, we show that the developed prior provides consistent selection of a fixed, finite true model, which is assumed to be in the model space. The true model is considered to be the smallest model that utilizes all of the covariates with non-zero coefficients and no covariates whose coefficient is zero. The assumption that the true model is finite is a more restrictive than the assumptions in Castillo et al. (2015) and corresponds to the framework of Johnson and Rossell (2012). Like Johnson and Rossell (2012), it is also assumed that p is increasing with n and that p is less than $n - p_0 - 1$ so that the design matrix can be assumed to be full rank. In Section 4, we describe a model search procedures that aims to provide a polynomial time solution to searching the model space. In lieu of random walks, the search is carried out in a semi-greedy fashion. A rooted tree is built out of the model space and vertices are visited in a particular order. We provide a few rules for pruning branches from the rooted tree. These rules allow for one to ignore the majority of the model space while finding good models in polynomial time. In Section 5, we compare different model space priors by analyzing the the diabetes data set from Efron et al. (2004). The design matrix is augmented with superfluous covariates that are correlated with the original covariates. The results of the data analysis show that the limiting Poisson prior provides good control of false positives while not sacrificing the discovery of true positives. In Section 6, we conclude the paper with a discussion of ongoing research extending

the analysis using the limiting Poisson prior and "objective" priors to obtain results similar to Castillo et al. (2015).

2. Construction of Model Space Priors

This section considers the motivation for various priors on the model space and proposes a new framework for constructing a prior on the model space. One important condition on the model space is that all models of the same size get the same prior probability. This leads to priors on the model space of the form

$$\pi_p(M_A) = \pi_p(|A|)\binom{p}{|A|}^{-1}. \tag{2}$$

The priors we consider will differ only in how they allocate prior probability to different model sizes — here represented by s — through $\pi_p(s)$. A prior satisfying (2) can be achieved through an infinite exchangeability assumption, though an assumption of finite exchangeability for each p provides the form in (2). Multiplicity correction priors are obtained by making an infinite exchangeability assumption and priors that penalize models based on their complexity can be obtained by making changes to the standard Beta-Binomial prior parameters. After these priors are reviewed, we provide a new formulation for constructing priors on the model space. In contrast to the infinite exchangeability assumption, this formulation does not view the model space through indicator variables. Rather, it views the model space as a partially ordered set and builds a prior by using the idea of "local" null hypotheses and alternatives, which leads naturally to a Poisson prior on the size of the model.

2.1. The Infinite Exchangeability Assumption

Each covariate $\mathbf{x}_i$ has an associated binary indicator γ_i, and model M_A can be equated with $\gamma_A = (\gamma_{A,1}, \ldots, \gamma_{A,p})$ where $\gamma_{A,i} = I(i \in A)$. The γ_i are assumed to form an infinitely exchangeable sequence of random variables. Such an assumption is reasonable when p is increasing and in the limit there are an infinite number of γ_i. This is one of the few instances where the distribution of the γ_i can be explicitly determined just from the infinite exchangeability assumption. By the de Finetti theorem, there exists a parameter $\phi \in (0, 1)$ and a mixing distribution $\pi(\phi)$ such that, when conditioned on ϕ, the γ_i are independent Bernoulli random variables with probability of success ϕ. The only difficulty is determining a good distribution for $\pi(\phi)$. The objective choice would be to use the Jeffreys prior, which is a Beta$(1/2, 1/2)$ distribution, while Scott and Berger (2010) argue in favor of the Beta$(1, 1)$ prior, which produces a uniform distribution over model dimension.

However, any fixed distribution can be subjectively chosen and it is common to use a Beta(a, b) prior. Note that the infinite exchangeability assumption requires that a and b are fixed.

Under the Beta(a, b) prior, the prior probability of the set of models with size s is

$$\pi_p(s) = \frac{\Gamma(a+b)\Gamma(p+1)\Gamma(s+a)\Gamma(p-s+b)}{\Gamma(b)\Gamma(a)\Gamma(s+1)\Gamma(p-s+1)\Gamma(p+a+b)}. \tag{3}$$

Fixing s and letting $p \to \infty$ provides

$$\pi_p(s) \to p^{-a} \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \frac{\Gamma(s+a)}{\Gamma(s+1)}.$$

Because of this, the Beta-Binomial priors only act to provide multiplicity correction through $\binom{p}{s}^{-1}$ in (2) and do not really provide for any further penalization for model complexity. Different values of a do provide for differences in the probability of sets of models of different dimensions, but these differences are essentially meaningless for helping to select models when p is increasing. In particular,

$$\frac{\pi_p(s)}{\pi_p(s-1)} = \left(\frac{s-1+a}{s}\right)\left(\frac{p-s+1}{p-s+b}\right), \quad \text{for } s = 1, \ldots, p. \tag{4}$$

This shows that the probability of the set of models with s covariates does not differ very much from the probability of the set of models with $s+1$ covariates. Such a lack of penalization makes it difficult to control false positives in regression. Moreover, the priors place probability one on the set of models where s increases as a fraction of p as p increases, which makes it impossible for the posterior probability of a finite true model to converge to one (Johnson and Rossell, 2012).

2.2. Complexity Penalizing Priors

In order to address issues with Beta-Binomial priors, Wilson et al. (2010) take $a = 1$ and $b = \lambda p$. This renders the prior expectation for s to be $1/\lambda$ and the limit of (4) to be $1/(\lambda+1)$ for fixed s as p increases. Thus, the prior for $\pi_p(s)$ behaves like a geometric distribution for s that are small with respect to p. This additional penalization provides some complexity control and the limiting geometric prior is a meaningful distribution on the non-negative integers.

A stronger assumption is made in Castillo et al. (2015), where it is assumed that

$$c_1 p^{-c_2} \leq \frac{\pi_p(s)}{\pi_p(s-1)} \leq c_3 p^{-c_4} \tag{5}$$

for positive c_i. A prior satisfying (5) can be achieved by taking $b = p^u$ for $u > 1$ in the Beta-Binomial prior. Assumption (5) causes the prior probability of $\pi_p(s)$ to place mass one on $s = 0$ as p increases to ∞, which is a meaningless distribution on the non-negative integers.

There is no justification for either of these approaches from first principles. Taking b to depend on p breaks the infinite exchangeability assumption and the finite exchangeability assumption for

each p could allow any possible distribution for $\pi_p(s)$. Moreover, though the assumption in (5) provides some nice theoretical results, at face value it seems to be much stronger than is necessary. When differentiating models with small s, the learning rates from the Bayes factors should be enough for a finite true model to beat competing models when there is a finite exchangeability assumption. Issues arises when the true model has a large dimension t (for instance, t is a fraction of p). Then the learning rates from the Bayes factors might break down and a stronger penalization of models with size larger than t is necessary. Thus, the requirement in (5) should only have to hold for large s and should be able to be ignored for small s. Moreover, the lack of a framework for interpreting (5) outside of its posterior properties seems problematic. The prior on the model space should represent a set of reasonable beliefs about competing models before the data is observed.

2.3. Prior Construction using Local Null hypotheses

Here we propose new criteria for constructing model space priors. These criteria utilize the idea that model M_A is the local hypothesis that $\beta_i = 0$ for $i \notin A$. The set of local alternatives over M_A are models that assume that $\beta_i \neq 0$ for at least one $i \notin A$. Define the descendant set of M_A to be the set of models $D(M_A) = \{M_B \in \mathcal{M} : A \subsetneq B\}$ and the children set of M_A to be $C(M_A) = \{B \in D(M_A) : |A| + 1 = |B|\}$. The two criteria criteria that we propose for constructing model space priors are

$$\pi(M_A) \propto \sum_{B \in C(M_A)} \pi(M_B) \tag{6}$$

$$\text{and} \quad \pi(M_A) \propto \sum_{B \in D(M_A)} \pi(M_B). \tag{7}$$

Before discussing the implications of (6) and (7), we discuss a natural condition on the model space when p increases to ∞.

Because each model can be represented by a sequence of indicator variables, the model space is equivalent to $G_p = \{0,1\}^{\times p}$. It is illuminating to consider what occurs as $p \to \infty$. Then $G_p \to G_\infty = \{0,1\}^{\times \mathbb{N}}$, which is the set of countable binary vectors. Suppose that, as investigators, we have a strong belief about whether a particular set of covariates should be included in the true model. For convenience assume that these covariates are labeled $i = 1, \ldots, \ell$. Instead of including these covariates in the set of tested covariates, we can simply declare the inclusion or exclusion of these variables in the base model and consider the reduced model space where the inclusion of these variable is not in question. This removal is equivalent to performing a left shift operation (denoted by L) on the G_∞ to form $L^\ell(G_\infty)$, which is isomorphic to G_∞. Now, we can force the condition that the renormalized prior on $L^\ell(G_\infty)$ is the same as the prior on G_∞. Let $\pi_\infty(s)$ be

the distribution over model sizes when $p \to \infty$ and a finite exchangeability assumption is made for each p. A simple limiting argument shows that the equality of the prior on G_∞ and $L^\ell(G_\infty)$ leads to the system of equations

$$\pi_\infty(s) = \frac{\pi_\infty(s+\ell)\frac{(s+\ell)!}{s!}}{\sum_{j=0}^{\infty}\pi_\infty(j+\ell)\frac{(j+\ell)!}{j!}} \tag{8}$$

for all integers ℓ and s that are non-negative. This system of equations is only satisfied by a Poisson distribution for $\pi_\infty(s)$.

The argument from the isometry in the infinite limit is somewhat dissatisfying, as it provides no means of determining a prior outside of the infinite case and needs to make a finite exchangeability assumption for each p. Fortunately, each of two simple constructions implied by (6) and (7) leads directly to a distribution for $\pi_p(M_A)$ that satisfies finite exchangeability for all p and provides a distribution for $\pi_p(s)$ that converges to a Poisson distribution as $p \to \infty$.

The first construction relates the a local null hypothesis M_A to the set of alternatives that are represented by the children of M_A, $C(M_A)$. In particular, it is assumed that

$$\pi_p(M_A) = \rho \sum_{M_{A'} \in C(M_A)} \pi_p(M_{A'}) \tag{9}$$

for some fixed $\rho > 0$. First, it is easy to confirm that this assumption forces a finite exchangeability condition for all p. Second, using the finite exchangeability condition and (9), we obtain

$$\pi_p(s)\binom{p}{s}^{-1} = \rho \times \pi_p(s+1)\binom{p}{s+1}^{-1}\binom{p-s}{1},$$

which provides $\pi_p(s) = \rho \times \pi_p(s+1) \times (s+1)$. This equality is only satisfied by a truncated Poisson distribution with rate $\lambda = \rho^{-1}$ that is restricted to $\{0, \ldots, p\}$.

The second construction replaces the relationship based on $C(M_A)$ with a relationship based on $D(M_A)$, the set of descendants of M_A. For this construction, it is assumed that

$$\pi_p(M_A) = \eta \sum_{M_{A'} \in D(M_A)} \pi_p(M_{A'}) \tag{10}$$

for some fixed $\eta > 0$. Once again, it is easy to confirm that this assumption forces a finite exchangeability condition for all p. Using the finite exchangeability condition and (10), we obtain

$$\pi_p(s)\binom{p}{s}^{-1} = \eta \times \sum_{\ell=1}^{p-s} \pi_p(s+\ell)\binom{p}{s+\ell}^{-1}\binom{p-s}{\ell},$$

which provides

$$\pi_p(s) = \eta \times \sum_{\ell=1}^{p-s} \pi_p(s+\ell)\binom{s+\ell}{\ell} \quad \text{for all} \quad s \in \{0, \ldots, p-1\}. \tag{11}$$

All that remains is to prove that the limiting distribution is a Poisson distribution with a rate that depends on η. Using the self-similarity property for model spaces with finite p and letting $p \to \infty$ provides an easy proof of such a fact.

THEOREM 1. *Suppose that for all $p > 0$, the distribution $\pi_p(s)$ for $s \in \{0, \ldots, p\}$ satistfies* (11). *Then*

$$\lim_{p\to\infty} \pi_p(s) = \frac{\lambda^s \exp(-\lambda)}{s!}$$

for $\lambda = \log(1 + \eta^{-1})$.

PROOF. It is easy confirm that $\pi_p(0) = \eta/(1+\eta)$ for all $p \in \mathbb{N}$. Write

$$\pi_p(s) = \frac{1}{1+\eta} \times \frac{N_p(s)}{Z_p} \quad \text{for} \quad s \in \{1, \ldots, p\}$$

with $N_p(p) = 1$ and $Z_1 = 1$. An induction argument shows that $N_p(s) = \eta\binom{p}{s} Z_{p-s}$ for $s \in \{1, \ldots, p-1\}$. Because $\pi_\infty(s)$ is non-zero for all s, the probability $\pi_\infty(1)$ is given by

$$\pi_\infty(1) = \frac{\eta}{1+\eta} \lim_{p\to\infty} \frac{pZ_{p-1}}{Z_p}$$

and pZ_{p-1}/Z_p must converge to a finite limit as p increases; call this limit λ. For $s > 1$, write

$$\frac{\binom{p}{s} Z_{p-s}}{Z_p} = \frac{1}{s!} \prod_{i=0}^{s-1} \frac{(p-i)Z_{p-i-1}}{Z_{p-i}}$$

In the limit as $p \to \infty$, each term in the product converges to λ. Thus,

$$\pi_\infty(s) = \frac{\eta}{1+\eta} \frac{\lambda^s}{s!},$$

which is the probability of s under a Poisson distribution with parameter $\lambda = \log(1 + \eta^{-1})$. □

Three remarks must be made about these results. First, whether considering the children or descendant condition, the limiting distribution over model sizes is Poisson. The only difference is the rate λ of the distribution. It is natural to require the descendant condition with $\eta = 1$, which gives $\lambda = \log(2) < 1$. However the same λ can be obtained from the children condition with $\rho = 1/\log(2)$. The subjective choice of which condition and what value of ρ or η to use allows investigators to change the behavior of the priors to reflect beliefs about the size of the finite true model.

Second, the Poisson construction is quite natural and comes from a simple reinterpretation of the model space. If one is working from first principles, the only arguments that make sense as p increases are infinite exchangeability and self-similarity of G_∞ and $L^\ell(G_\infty)$. The only priors that come from these assumptions are the multiplicity correction prior (like the Beta-Binomial with

fixed parameters) or the limiting Poisson distribution. The former is known to behave poorly, while Section 3 will show that the latter provides nice consistency results. Moreover, the children and descendant conditions provide an interpretable construction of model space priors from first principles even when p is fixed, which is impossible under an infinite exchangeability assumption.

Third, it is useful to compute $\pi_\infty(s)/\pi_\infty(s-1)$ for $s > 0$ and compare the result to (5). In particular, the ratio $\pi_\infty(s)/\pi_\infty(s-1)$ is simply λs^{-1}. This is a weaker condition than (5), which forces $\pi_p(s)/\pi_p(s-1) \to 0$ as $p \to \infty$. In terms of finite s, it is not clear why the latter is at all desirable. It does produce an assistance to the Bayes factor from local priors when $p \propto n$ that similar to that obtained using non-local priors for the regression coefficients. In contrast, the Poisson prior maintains a reasonable penalty for finite s and has that penalty adaptively increase as s increases. In fact, when $s \propto p^u$ for $0 < u < 1$, then the probabilities for the Poisson satisfy (5) with all $c_1 = c_3 = \lambda$ and $c_2 = c_4 = u$. The fact that the limiting distribution from our constructions provides an adaptive form of (5) is quite surprising. This suggest that the Poisson prior can provide not only consistency results for finite true models, but that all of the results of Castillo et al. (2015) should be able to be reconstructed using it.

3. Posterior Consistency

In this section, we prove posterior consistency using the descendant and children priors when the true model is assumed to be $M_T \in \mathcal{M}$ with T being fixed and $|T|$ being finite. The true values of the parameters of M_T are given by $(\boldsymbol{\beta}_{0,T}, \boldsymbol{\beta}_T, \tau_T)$. For ease of exposition, the prior on the regression parameters is assumed to be the intrinsic prior (Berger and Pericchi, 1996; Casella et al., 2009; Moreno et al., 2010; Girón et al., 2010; Casella and Moreno, 2012; Womack et al., 2014). However, all of the consistency results in this section hold for large families of scaled mixtures of g-priors. The only differences are the details in the rates of convergence of Bayes factors for models versus the base model.

3.1. Bayes Factors

The model corresponding to $A = \emptyset$ is denoted by M_0 and it is assumed that M_0 is fixed. For generic $A \subset \{1, \ldots, p\}$, the Bayes factor for model M_A versus M_0 is given by

$$BF_{A,0} = \int \left(\frac{n + (|A| + p_0 + 1)w}{(1 - R_A^2)n + (|A| + p_0 + 1)w} \right)^{\frac{n-p_0}{2}} \times \left(\frac{(|A| + p_0 + 1)w}{n + (|A| + p_0 + 1)w} \right)^{\frac{|A|}{2}} \pi(w) \dot{w} \tag{12}$$

where

$$R_A^2 = \frac{\mathbf{y}'(\mathbf{H}_A - \mathbf{H}_0)\mathbf{y}}{\mathbf{y}'(\mathbf{I} - \mathbf{H}_0)\mathbf{y}}, \tag{13}$$

is the coefficient of determination for model M_A, $\mathbf{H}_0$ is the hat matrix associated with $\mathbf{X}_0$, and $\mathbf{H}_A$ is the hat matrix associated with $(\mathbf{X}_0\ \mathbf{X}_A)$. The parameter w in (12) is an additional scaling of the variance in a scaled mixture of g-priors (Womack et al., 2014). For the intrinsic prior, the prior density for w is given by

$$\pi(w) = \frac{I(0 < w < 1)}{\pi w^{1/2}(1-w)^{1/2}}.$$

The Zellner-Siow prior takes the prior for w to be Gamma$(1/2, b/2)$ with $b = 1$. The hyper parameter b can be given an additional Gamma prior, which produces a beta distribution of the second kind for w. Some of the details of the proof we present here use the fact that the prior for w is only supported on the unit interval for the intrinsic prior.

3.1.1. Decomposition of the coefficient of determination

For $A \not\supseteq T$, $1 - R_A^2$ in (12) can be written as

$$\begin{aligned}
1 - R_A^2 &= \frac{\xi_1(A) + \xi_2(A)}{\xi_1(A) + \xi_2(A) + \xi_3(A) + \xi_4(A)}, \\
\text{where} \quad \xi_1(A) &= \tau_T \mathbf{y}'(\mathbf{I} - \mathbf{H}_{A\cup T})\mathbf{y}, \\
\xi_2(A) &= \tau_T \mathbf{y}'(\mathbf{H}_{A\cup T} - \mathbf{H}_A)\mathbf{y}, \\
\xi_3(A) &= \tau_T \mathbf{y}'(\mathbf{H}_A - \mathbf{H}_{A\cap T})\mathbf{y}, \\
\text{and} \quad \xi_4(A) &= \tau_T \mathbf{y}'(\mathbf{H}_{A\cap T} - \mathbf{H}_0)\mathbf{y}.
\end{aligned}$$

The four ξ_j are independent non-central χ^2 random variables following the distributions

$$\begin{aligned}
\xi_1(A) &\sim \chi^2_{n-p_0-|A\cup T|}(0) \\
\xi_2(A) &\sim \chi^2_{|A\cup T|-|A|}(n\delta_2(A)) \\
\xi_3(A) &\sim \chi^2_{|A|-|A\cap T|}(n\delta_3(A)) \\
\xi_4(A) &\sim \chi^2_{|A\cap T|}(n\delta_4(A))
\end{aligned}$$

with

$$\begin{aligned}
\delta_2(A) &= \lim_{n\to\infty} \frac{\tau_T(\boldsymbol{\beta}'_{0,T}\ \boldsymbol{\beta}'_T)(\mathbf{X}_0\ \mathbf{X}_T)'(\mathbf{H}_{A\cup T} - \mathbf{H}_A)(\mathbf{X}_0\ \mathbf{X}_T)(\boldsymbol{\beta}'_{0,T}\ \boldsymbol{\beta}'_T)'}{n} \\
\delta_3(A) &= \lim_{n\to\infty} \frac{\tau_T(\boldsymbol{\beta}'_{0,T}\ \boldsymbol{\beta}'_T)(\mathbf{X}_0\ \mathbf{X}_T)'(\mathbf{H}_A - \mathbf{H}_{A\cap T})(\mathbf{X}_0\ \mathbf{X}_T)(\boldsymbol{\beta}_{0,T}\ \boldsymbol{\beta}_T)}{n} \\
\delta_4(A) &= \lim_{n\to\infty} \frac{\tau_T(\boldsymbol{\beta}'_{0,T}\ \boldsymbol{\beta}'_T)(\mathbf{X}_0\ \mathbf{X}_T)'(\mathbf{H}_{A\cap T} - \mathbf{H}_0)(\mathbf{X}_0\ \mathbf{X}_T)(\boldsymbol{\beta}'_{0,T}\ \boldsymbol{\beta}'_T)'}{n}.
\end{aligned}$$

In general, for $B \subsetneq B' \subseteq \{1, \ldots, p\}$, the difference in hat matrices $\mathbf{H}_{B'} - \mathbf{H}_B$ is given by

$$\mathbf{H}_{B'} - \mathbf{H}_B = (\mathbf{I} - \mathbf{H}_B)\mathbf{X}_{B' \cap B^c} \left(\mathbf{X}'_{B' \cap B^c}(\mathbf{I} - \mathbf{H}_B)\mathbf{X}_{B' \cap B^c}\right)^{-1} \mathbf{X}'_{B' \cap B^c}(\mathbf{I} - \mathbf{H}_B),$$

which we will denote by $\mathbf{H}_{B' \perp B}$. Thus, the non-centrality parameters can be written as

$$\begin{aligned}
\delta_2(A) &= \lim_{n\to\infty} \frac{\tau_T \boldsymbol{\beta}'_{T\cap A^c} \mathbf{X}'_{T\cap A^c} \left(\mathbf{H}_{(A\cup T)\perp A}\right) \mathbf{X}_{T\cap A^c} \boldsymbol{\beta}_{T\cap A^c}}{n} \\
&= \lim_{n\to\infty} \frac{\tau_T \boldsymbol{\beta}'_{T\cap A^c} \mathbf{X}'_{T\cap A^c} \left(\mathbf{I} - \mathbf{H}_A\right) \mathbf{X}_{T\cap A^c} \boldsymbol{\beta}_{T\cap A^c}}{n} \\
\delta_3(A) &= \lim_{n\to\infty} \frac{\tau_T \boldsymbol{\beta}'_{T\cap A^c} \mathbf{X}'_{T\cap A^c} \left(\mathbf{H}_{A\perp(A\cap T)}\right) \mathbf{X}_{T\cap A^c} \boldsymbol{\beta}_{T\cap A^c}}{n} \\
\delta_4(A) &= \lim_{n\to\infty} \frac{\tau_T \boldsymbol{\beta}'_{T\cap A} \mathbf{X}'_{T\cap A} \left(\mathbf{H}_{(A\cap T)\perp 0}\right) \mathbf{X}_{T\cap A} \boldsymbol{\beta}_{T\cap A}}{n} \\
&= \lim_{n\to\infty} \frac{\tau_T \boldsymbol{\beta}'_{T\cap A} \mathbf{X}'_{T\cap A} \left(\mathbf{I} - \mathbf{H}_0\right) \mathbf{X}_{T\cap A} \boldsymbol{\beta}_{T\cap A}}{n}.
\end{aligned}$$

The non-centrality parameters satisfy $\delta_i(A) \geq 0$ for $i = 2, \ldots, 4$ with $\sum_{i=2}^4 \delta_i(A) = \delta_4(T) > 0$. Further, $\delta_4(A) = 0$ if and only if $A \cap T = \emptyset$, $\delta_2(A) = 0$ if and only if $\mathbf{X}'_{A^c\cap T}(\mathbf{I} - \mathbf{H}_A)/n \to \mathbf{0}$ as $n \to \infty$, and $\delta_3(A) = 0$ if and only if $\mathbf{X}'_{A^c\cap T}(\mathbf{I} - \mathbf{H}_{A\cap T})\mathbf{X}_{A\cap T^c}/n \to \mathbf{0}$ as $n \to \infty$. If $A \supsetneq T$ then $\xi_2(A) = \delta_2(A) = \delta_3(A) = 0$, and if $A \subsetneq T$ then $\xi_3(A) = \delta_3(A) = 0$.

Finally, we note that any model M_A provides a novel decomposition of the true expectation of $\mathbf{y}$ into orthogonal pieces; $E[\mathbf{y}] = \sum_{i=0}^4 \boldsymbol{\mu}_i$ where

$$\begin{aligned}
\boldsymbol{\mu}_1(A) &= (\mathbf{I} - \mathbf{H}_{A\cup T}) \mathbf{X}_T \boldsymbol{\beta}_T = \mathbf{0}, \\
\boldsymbol{\mu}_2(A) &= \left(\mathbf{H}_{(A\cup T)\perp A}\right) \mathbf{X}_{A^c\cap T} \boldsymbol{\beta}_{A^c\cap T}, \\
\boldsymbol{\mu}_3(A) &= \left(\mathbf{H}_{A\perp(A\cap T)}\right) \mathbf{X}_{A^c\cap T} \boldsymbol{\beta}_{A^c\cap T}, \\
\boldsymbol{\mu}_4(A) &= \left(\mathbf{H}_{(A\cap T)\perp 0}\right) \mathbf{X}_{A\cap T} \boldsymbol{\beta}_{A\cap T}, \\
\text{and} \quad \boldsymbol{\mu}_0(A) &= \mathbf{X}_0 \boldsymbol{\beta}_0.
\end{aligned}$$

The non-centrality parameter δ_i is given by the limit as $n \to \infty$ of $\boldsymbol{\mu}'_i \boldsymbol{\mu}_i / n$.

3.1.2. Bayes Factor for finite models

The fixed finite model, M_A, is assumed to have $|A| + p_0$ total regressors. Taking the limit as $n \to \infty$, the Bayes factor for M_A versus M_T is given by

$$BF_{A,T} \to \left(\frac{\xi_1(T)}{\xi_1(A) + \xi_2(A)}\right)^{\frac{n-p_0}{2}} \frac{\left(\frac{|A|+p_0+1}{n}\right)^{\frac{|A|}{2}} I(A)}{\left(\frac{|T|+p_0+1}{n}\right)^{\frac{|T|}{2}} I(T)} \tag{14}$$

where

$$I(B) = \int \exp\left(-\frac{(\delta_3(B) + \delta_4(B))}{(1 + \delta_2(B))}\frac{(|B| + p_0 + 1)w}{2}\right) w^{\frac{|B|}{2}} \pi(w) \text{w}$$

for fixed B. If $\delta_2(A) = 0$, then the first term in (14) converges to

$$\left(\frac{\xi_1(T)}{\xi_1(A) + \xi_2(A)}\right)^{\frac{n-p_0}{2}} \to \exp\left(-\frac{\xi_2(A)}{2}\right)$$

as $n \to \infty$. When $\delta_2(A) > 0$, then this term converges to

$$\left(\frac{\xi_1(T)}{\xi_1(A) + \xi_2(A)}\right)^{\frac{n-p_0}{2}} \to \exp\left(-\frac{n - p_0}{2}\log\left(1 + \delta_2(A)\right)\right).$$

These limiting Bayes factors provide the usual exponential and power learning rates for M_T over M_A for finite models.

3.1.3. Requirements on the design matrix

It is now necessary to make some comments about conditions on $(\mathbf{X}_0\ \mathbf{X})$ to ensure that the χ^2 distributions in the Bayes factors behave as expected. The first condition is that $p \leq n - p_0 - 1$ and that

$$\lim_{n\to\infty} \frac{(\mathbf{X}_0\ \mathbf{X})'(\mathbf{X}_0\ \mathbf{X})}{n}$$

converges to a positive definite operator so that the Bayes factor from the intrinsic prior is well-defined. As in Johnson and Rossell (2012), we let $\lambda_{1,p} \leq \cdots \leq \lambda_{p+p_0,p}$ be the ordered eigenvalues of $(\mathbf{X}_0\ \mathbf{X})'(\mathbf{X}_0\ \mathbf{X})$ and assume that there are constants $0 < \lambda_0 \leq \lambda_\infty < \infty$ with $n\lambda_0 \leq \lambda_i \leq n\lambda_\infty$ for $i = 1, \ldots, p+p_0$ for all $p \geq 0$. An additional important assumption is that there are no observations with undue asymptotic influence, which is described as a condition on the the ℓ^2 norm of the vector of regressors for any individual (Knight and Fu, 2000). In this setting, it is assumed that for any fixed finite set of indices, say $a_1, \ldots, a_j$, no observation has too much influence. Specifically

$$\lim_{n\to\infty} \frac{\max_{1\leq i\leq n}\left(x_{0,i,1}^2 + \cdots + x_{0,i,p_0}^2 + x_{i,a_1}^2 + \cdots + x_{i,a_j}^2\right)}{n} = 0$$

These conditions ensure that for all models $A \not\supseteq T$ the non-centrality parameters $\delta_2(A)$ are uniformly bounded away from zero. Let δ_2^* be a positive lower bound for $\{\delta_2(A)\ :\ A \not\supseteq T\}$. These assumptions also ensure that for the set of descendant models of the true model, $D(M_T)$, the set of χ^2 random variables defined by $\xi_3(A)$ for $M_A \in D(M_T)$ have a reasonable joint distribution and do not accumulate mass (or degenerate) at any one particular value as n, and thus p, increases. All manner of other assumptions, such as those based on compatibility numbers or mutual coherence

(Castillo et al., 2015), are designed to ensure some nice set of properties for the non-centrality parameters and χ^2 random variables. The assumptions made here are more restrictive than necessary, but are sufficient to provide asymptotic consistency when M_T is a finite model.

3.2. Proof of Posterior Concentration on M_T

In this section, we prove state and prove the main theorem of the paper. I the Poisson prior is the limiting prior for $\pi_p(s)$ and the true model M_T is fixed and finite, then the posterior probability of M_T converges to one as $n \to \infty$. The proof relies on the Bayes factor learning rates and the strong penalization induced by the prior distribution. Before stating and proving the theorem formally, we provide some heuristic arguments for its validity. f

First, for a model M_A that does not nest M_T, the Bayes factor learning rate for M_T over M_A is exponential. For a given model size $s \geq |T|$, there are exactly

$$\sum_{j=0}^{|T|-1} \binom{|T|}{j}\binom{p-|T|}{s-j} = \binom{p}{s} - \binom{p-|T|}{s-|T|}$$

models that do not nest M_T (when $|T| = 0$, the sum is taken to be empty and thus equal to zero). The asymptotic prior odds of this set of models versus M_T is

$$\begin{aligned}\frac{\lambda^{s-|T|}(|T|)!}{s!}\binom{p}{|T|}\binom{p}{s}^{-1}&\left[\binom{p}{s} - \binom{p-|T|}{s-|T|}\right] \\ &= \frac{\lambda^{s-|T|}(|T|)!}{s!}\left[\binom{p}{|T|} - \binom{s}{|T|}\right] \\ &\leq \frac{\lambda^s}{s!}\left(\frac{p}{\lambda}\right)^{|T|}.\end{aligned}$$

In this case, the exponential learning M_T versus M_A provides sufficient penalization to overcome $p^{|T|}$ and the sum of the posterior odds of these models converges to zero as n increases. The additional difficulty in proving convergence for the set of all models not nesting M_T comes from models whose dimension is essentially p.

Second, for a model M_A that nests M_T, the Bayes factor learning rate is for M_T over M_A is $n^{(|A|-|T|)/2}$. For a given model size $s > |T|$, there are exactly $\binom{p-|T|}{s-|T|}$ models that nest M_T. This provides an asymptotic prior odds of $\lambda^{s-|T|}/(s-|T|)!$ for this set of models versus M_T. Summing $(n^{-1/2}\lambda)^{s-|T|}/(s-|T|)!$ over all s provides a summable series that converges to zero as $n \to \infty$. The additional difficulty of proving convergence for the sum of all models that nest M_T comes from infinite models whose dimension is essentially p, where $1 - R^2$ converge to 0 and the learning rate from the Bayes factor can break down. Here, the properties of the intrinsic prior are quite useful,

though the proof of convergence does follow for other scaled mixtures of g-priors. We proceed to state and prove Theorem 2.

THEOREM 2. *Let $M_T \in \mathcal{M}$ be a fixed finite true model and assume the model space $\mathcal{M}$, with $p \leq n - p_0 - 1$, satisfies the following*

(a) the eigenvalues of $(\mathbf{X}_0\ \mathbf{X})'(\mathbf{X}_0\ \mathbf{X})/n$ are uniformly bounded above and uniformly bounded away from zero for all n,

(b) for each finite model, each observation has negligible influence as $n \to \infty$,

(c) the limiting prior on model complexity, $\pi_\infty(s)$, is a Poisson distribution with parameter λ,

(d) and, the prior on regression parameters is the intrinsic prior.

Then posterior probability of M_T converges to one almost surely.

PROOF. The posterior probability of the true model M_T will converge to one if

$$\sum_{A \neq T} PO_{A,T} \to 0 \quad \text{almost surely,} \tag{15}$$

where

$$PO_{A,T} = \frac{BF_{A,0}\pi_p(M_A)}{BF_{T,0}\pi_p(M_T)}$$

is the posterior odds of M_A versus M_T. Taking the limit as $n \to \infty$, the prior odds converge to

$$\frac{\pi_p(M_A)}{\pi_p(M_T)} = \frac{\pi_p(|A|)(|A|)!(p-|A|)!}{\pi_p(|T|)(|T|)!(p-|T|)!} \to \lambda^{|A|-|T|}\frac{(p-|A|)!}{(p-|T|)!}.$$

The quantity in (15) decomposes into two parts

$$S_{\not\supseteq}(T) = \sum_{A \not\supseteq T} PO_{A,T} \tag{16}$$

$$S_{\supsetneq}(T) = \sum_{A \supsetneq T} PO_{A,T}, \tag{17}$$

where (16) sums the posterior odds over all models that do not next M_T and (17) sums over all models that nest M_T and are not M_T.

The proof that the posterior probability of M_T converges to one is contained in the following two lemmas. The first shows that (16) converges to zero and the proof is not too complicated while the second shows that (17) converges to zeros and more care is needed. In the proofs of the lemmas, the limiting Poisson prior is substituted for the prior induced by the descendant or children condition for each finite p. □

LEMMA 1. *Assume the conditions of Theorem 2, then $S_{\not\supseteq}(T)$ in* (16) *converges to zero almost surely.*

PROOF. Recall that conditions *1* and *2* in the statement of the theorem are enough to ensure that there exists $\delta_2^* > 0$ such that $\delta_2(A) > \delta_2^*$ for all $A \not\supseteq T$. In the case where δ_2^* can be taken to be greater than one, the proof that $S_{\not\supseteq}(T)$ converges to zero is relatively easy. With probability one, the Bayes factor $BF_{A,0}$ for each model is asymptotically bounded above (quite crudely) by

$$BF_{A,0} \leq \left(\frac{1+\delta_4(T)}{\delta_2^*}\right)^{\frac{n-p_0}{2}} \left(\frac{1}{2}\right)^{\frac{|A|}{2}}$$

and so the posterior odds of models of size s versus M_T is asymptotically bounded above by

$$\frac{\lambda^s}{s!(\sqrt{2})^s} \left(\frac{p\sqrt{n}}{\lambda\sqrt{|T|+p_0+1}}\right)^{|T|} \frac{1}{I(T)} \left(\frac{1}{\delta_2^*}\right)^{\frac{n-p_0}{2}}.$$

Summing this from $s=0$ to $s=\infty$ provides

$$S_{\not\supseteq}(T) \leq \exp\left(\frac{\lambda}{\sqrt{2}}\right) \left(\frac{p\sqrt{n}}{\lambda\sqrt{|T|+p_0+1}}\right)^{|T|} \frac{1}{I(T)} \left(\frac{1}{\delta_2^*}\right)^{\frac{n-p_0}{2}},$$

which converges to zero as $n \to \infty$ because $p \leq n$ and $|T|$ is fixed.

For the proof when δ_2^* is necessarily less than one, we have to break the sum into two pieces. The first piece looks at models whose dimension is less than $K_\alpha = \lfloor n\alpha\delta_2^* \rfloor$ for some $\alpha \in (0,1)$. The limiting value of the $1 - R_A^2$ for these models is bounded below by $(1+(1-\alpha)\delta_2^*)/(1+\delta_4(T))$. The proof that the sum of the posterior odds for models that do not nest M_T and whose dimensions are less than K_α proceeds using the same simple argument as the case when $\delta_2^* > 1$.

All that remains is to show that the sum of the posterior odds for models whose dimensions are at least $K_\alpha + 1$ and that do not nest M_T converges to zero. When $|A|/n \to r_A > \alpha\delta_2^* > 0$, the integral for the Bayes factor given in (12) is asymptotically bounded above by

$$\begin{aligned} BF_{A,0} &\leq \left(\frac{1+\delta_4(T)}{1-r_A+\delta_2(A)}\right)^{\frac{n-p_0}{2}} \left(\frac{r_A}{1+r_A}\right)^{\frac{r_A n}{2}} \\ &\leq \left(\frac{1+\delta_4(T)}{\delta_2^*}\right)^{\frac{n-p_0}{2}} \left(\frac{r_A}{1+r_A}\right)^{\frac{r_A n}{2}}. \end{aligned}$$

Using the Stirling approximation for $s!$, the posterior odds of models of size s with $s/n \to r_s > \alpha\delta_2^* > 0$ is asymptotically bounded above by

$$\begin{aligned} &\frac{\lambda^{r_s n}\exp(r_s n)}{\sqrt{2\pi r_s n}(r_s n)^{r_s n}} \frac{p^{|T|}n^{|T|/2}}{\lambda^{|T|}(|T|+p_0+1)^{|T|/2}I(T)} (\delta_2^*)^{-\frac{n-p_0}{2}} \left(\frac{r_s}{1+r_s}\right)^{\frac{r_s n}{2}} \\ =&C\left(\frac{\lambda^{r_s}e^{r_s}}{\sqrt{\delta_2^*}\sqrt{r_s^{r_s}(1+r_s)^{r_s}}}\right)^n \frac{(p\sqrt{n})^{|T|}}{\sqrt{r_s n}\ n^{r_s n}} \\ \leq&C\left(\frac{\max\{\lambda^{\alpha\delta_2^*},1\}e^1}{\sqrt{\delta_2^*}\sqrt{(\alpha\delta_2^*)(1+\alpha\delta_2^*)^{\alpha\delta_2^*}}}\right)^n \frac{(p\sqrt{n})^{|T|}}{\sqrt{\alpha\delta_2^*}\sqrt{n}\ n^{r_s n}} \end{aligned}$$

for some constant C that does not depend on r. The sum for $s > K_\alpha$ can be approximated by the integral

$$\begin{aligned}
&\int_{\alpha\delta_2^*}^1 C\left(\frac{\max\{\lambda^{\alpha\delta_2^*},1\}e^1}{\sqrt{\delta_2^*}\sqrt{(\alpha\delta_2^*)(1+\alpha\delta_2^*)^{\alpha\delta_2^*}}}\right)^n \frac{(p\sqrt{n})^{|T|}p}{\sqrt{\alpha\delta_2^*}\sqrt{n}\ n^{rn}}\mathrm{r}\\
=&C\left(\frac{\max\{\lambda^{\alpha\delta_2^*},1\}e^1}{\sqrt{\delta_2^*}\sqrt{(\alpha\delta_2^*)(1+\alpha\delta_2^*)^{\alpha\delta_2^*}}}\right)^n \frac{(p\sqrt{n})^{|T|}p}{\sqrt{\alpha\delta_2^*}\sqrt{n}}\frac{1}{n\log(n)}\left(\frac{1}{n^{\alpha\delta_2^* n}}-\frac{1}{n^n}\right)\\
\leq&C\left(\frac{\max\{\lambda^{\alpha\delta_2^*},1\}e^1}{\sqrt{\delta_2^*}\sqrt{(\alpha\delta_2^*)(1+\alpha\delta_2^*)^{\alpha\delta_2^*}}}\right)^n \frac{(p\sqrt{n})^{|T|}p}{\sqrt{\alpha\delta_2^*}\sqrt{n}}\frac{1}{n\log(n)}\left(\frac{1}{n^{\alpha\delta_2^* n}}\right),
\end{aligned}$$

which converges to zero due to the super-exponential decay of $n^{-\alpha\delta_2^* n}$. □

LEMMA 2. *Assume the conditions of Theorem 2, then* $S_{\supsetneq}(T)$ *in* (17) *converges to zero almost surely.*

PROOF. Let $\mathcal{M}_s(T)$ be the set of models $\{M_A \in \mathcal{M} : |A| = k \text{ and } A \supsetneq T\}$. The sum of the posterior odds of the descendant set $D(M_A)$ is

$$S_{\supsetneq}(T) = \sum_{s=|T|+1}^{p} \frac{\lambda^{s-|T|}}{(s-|T|)!} BF(\mathcal{M}_s(T)) = \sum_{s=|T|+1}^{p} PO(\mathcal{M}_s(T))$$

where

$$BF(\mathcal{M}_s(T)) = \binom{p-|T|}{s-|T|}^{-1} \sum_{A\in\mathcal{M}_s(T)} BF_{A,T}$$

is the average Bayes factor over the set $\mathcal{M}_s(T)$.

First, we consider the case when $s/n \to 0$. Each Bayes' factor is bounded above by

$$\begin{aligned}
BF_{A,T} &\leq \frac{1}{I(T)}\left(\frac{n}{|T|+p_0+1}\right)^{\frac{|T|}{2}}\left(\frac{|A|+p_0+1}{n+|A|+p_0+1}\right)^{\frac{|A|}{2}}\left(1+\frac{\xi_3(A)}{\xi_1(A)}\right)^{\frac{n-p_0}{2}}\\
&= \frac{(|A|+p_0+1)^{\frac{|T|}{2}}}{I(T)(|T|+p_0+1)^{\frac{|T|}{2}}}\left(\frac{|A|+p_0+1}{n+|A|+p_0+1}\right)^{\frac{|A|-|T|}{2}}\left(1-\frac{\xi_3(A)}{\xi_1(A)+\xi_3(A)}\right)^{-\frac{n-p_0}{2}}\\
&\leq \frac{(|A|+p_0+1)^{\frac{|T|}{2}}}{I(T)(|T|+p_0+1)^{\frac{|T|}{2}}}\left(\frac{|A|+p_0+1}{n+|A|+p_0+1}\right)^{\frac{|A|-|T|}{2}}\exp\left(\frac{(n-p_0)\xi_3(A)}{2\xi_1(A)}\right).
\end{aligned}$$

Define the random variables

$$W_n(A) = \exp\left(\frac{(n-p_0)\xi_3(A)}{2\xi_1(A)}\right)$$

where the $\xi_3(A)$ follow χ^2 distributions with degrees of freedom $|A|-|T|$. Because $|A|/n \to \infty$, each $W_n(A)$ converges in distribution to $W_\infty(A)$ where

$$W_\infty(A) = \exp\left(\frac{\xi_3(A)}{2}\right).$$

The $W_\infty(A)$ do not have finite expectation for all $M_A \in \mathcal{M}_s(T)$ for all $s > |T|$. However, though the random variables are not independent, each is in the domain of attraction of a stable law with index one. Conditions *(a)* and *(b)* ensure that in the limit the average for models of size s

$$\overline{W}_n(s) = \binom{p-|T|}{p-s}^{-1} \sum_{M_A \in \mathcal{M}_s(T)} W_n(A)$$

converges in distribution to a random variable $W^*(s)$ that is a stable distribution with stability index one. Conditions *(a)* and *(b)* further provide that the collection $W^*(s)^{\frac{1}{s-|T|}}$ satisfy the following. For any $\epsilon > 0$ there exists a $\Lambda > 0$ such that

$$P\left(W^*(s)^{\frac{1}{s-|T|}} > \Lambda\right) < \epsilon.$$

for all $s > |T|$.

With ϵ and Λ as above, the posterior odds for $\mathcal{M}_s(T)$ is bounded above by

$$PO(\mathcal{M}_s(T)) \leq C\left(\frac{\lambda\Lambda}{\sqrt{n}}\right)^{s-|T|}$$

with probability at least $1-\epsilon$ for all $s > |T|$, where C is a constant that depends on T but is independent of s. Taking the sum over $s - |T| > 0$ provides that

$$\sum_{s\ :s/n\to 0} PO(\mathcal{M}_s(T)) \leq C\frac{\lambda\Lambda}{\sqrt{n}-\lambda\Lambda}$$

with probability at least $1-\epsilon$ for n sufficiently large. Thus

$$\sum_{s\ :s/n\to 0} PO(\mathcal{M}_s(T)) \to 0 \quad \text{almost surely.}$$

Second, we consider the case when $s/n \to r_s \in (0,1)$. A more subtle result about the Bayes factor is necessary for the sum over such models to converge (see Moreno et al. (2010)). We have $1 - R_A^2 \to (1-r_A)/(1+\delta_4(T))$ with $r_A \in (0,1)$ and that the integral in (12) is less than the maximum of the integrand, which is achieved when

$$\frac{(|A|+p_0+1)}{n}\sin^2\left(\frac{\pi\theta}{2}\right) = \min\left(1, \frac{|A|(1-R_A^2)}{(n-p_0)R_A^2 - |A|}\right).$$

As $n \to \infty$, this maximum is given by $\min(1, r_A\delta_4(T)^{-1})$. If $0 < r_A \leq \delta_4(T)$, then

$$BF_{A,0} \leq (1+\delta_4(T))^{\frac{n-p_0}{2}}\left(\frac{\delta_4(T)}{\delta_4(T)+r_A}\right)^{\frac{r_A n}{2}}.$$

In the case where $0 < \delta_4(T) \leq r_A < 1$, the inequality becomes

$$BF_{A,0} \leq 2^{\frac{p_0}{2}}(1+\delta_4(T))^{\frac{n-p_0}{2}}\left(\frac{2^{r_A}}{2+\delta_4(T)-r_A}\right)^{\frac{n}{2}}.$$

Consider the sum of the posterior odds when $0 < \frac{s}{n} \leq \min(\delta_4(T), 1)$ and $s/n \to r_s > 0$. Then, for a particular s,

$$PO(\mathcal{M}_{s+|T|}(T)) \leq \left(\frac{n}{|T|+p_0+1}\right)^{\frac{|T|}{2}} \frac{1}{I_T}\frac{\lambda^s}{s!}\left(\frac{\delta_4(T)}{\delta_4(T)+r_s}\right)^{\frac{s}{2}}$$
$$\lessapprox \left(\frac{n}{|T|+p_0+1}\right)^{\frac{|T|}{2}} \frac{1}{I_T}\frac{1}{\sqrt{2\pi r_s n}}\left(\frac{\lambda e^1 \delta_4(T)^{1/2}}{r_s n(\delta_4(T)+r_s)^{1/2}}\right)^{r_s n}.$$

The sum over all such s can be approximated by the integral

$$\left(\frac{n}{|T|+p_0+1}\right)^{\frac{|T|}{2}} \frac{n^{1/2}}{I_T\sqrt{2\pi}}\int_0^{\min(1,\delta_4(T))} \frac{1}{\sqrt{r}}\left(\frac{\lambda e^1 \delta_4(T)^{1/2}}{rn(\delta_4(T)+r)^{1/2}}\right)^{rn} \dot{\mathrm{r}},$$

which converges to zero at a super-exponential rate as n increases.

If $\delta_4(T) < 1$, then the sum of the posterior odds over the models whose size s is a least $n\delta_4(T)$ with $s/n \to r_s < 1$ can be approximated by the integral

$$2^{p_0/2}\left(\frac{n}{|T|+p_0+1}\right)^{\frac{|T|}{2}} \frac{n^{1/2}}{I_T\sqrt{2\pi}}\int_{\delta_4(T)}^1 \frac{1}{\sqrt{r}}\left(\frac{2\lambda e}{rn(2+\delta_4(T)-r)^{1/r}}\right)^{rn} \dot{\mathrm{r}},$$

which also converges to zero at a super-exponential rate as n increases.

The final part of the proof concerns those models M_A with $r_A = 1$. These are models where $1 - R_A^2 \to 0$ and can be the most problematic for the Bayes factor. The Bayes' factor for a model to the base model is bounded above by

$$BF_{A,0} \lessapprox \left(\frac{n}{|T|+p_0+1}\right)^{\frac{|T|}{2}} \frac{(1+\delta_4(T))^{-\frac{n-p_0}{2}}}{I_T} \times 2^{n-|A|}\exp\left(-\frac{\psi_1(A)}{2(1+\delta_4(T))}\right),$$

The proof is finished using similar arguments to those for the case $|A|/n \to 0$. The differences are as follows: first, here $(n-|A|)/n \to 0$; second, the distributions $\exp\left(-\frac{\psi_1(A)}{2(1+\delta_4(T))}\right)$ has moments of all positive orders; and third, the prior odds of each model converges to zero at a super-exponential rate. □

4. Computational Algorithm

One major drawback of performing discrete model selection (or comparison) when p is large is the exponential growth of the model space with p. In order to find the model the highest posterior probability, we must perform an exhaustive search of the model space, which is impossible. However, the local null hypothesis nature of each model suggests a method for searching through the model space it in polynomial time. The limiting Poisson prior guarantees that the posterior probability of the true model converges to one as n increases. This, coupled with the strong penalization of models with large dimension allows us to essentially ignore models whose size is too large. Models whose

sizes are even only slightly greater than t, the size of the true model M_T, will get small posterior probabilities and very large models relative to t will receive negligible posterior probability.

4.1. A Greedy Ordering of Models

In order to define a search algorithm, we embellish the definition of a model and create a partial ordering on the set of models that can be computed as it is explored. Associate to each model $M_A \in \mathcal{M}$ a non-negative real number $Ev(M_A)$, which represents the evidence in favor of M_A. We will always take $Ev(M_A) = PO_{A,0}$, the posterior odds of M_A to M_0. For non-empty $A \subset \{1, \ldots, p\}$ order the elements of A in the following fashion

1. Let $a_{(1)} = \min\{a : a \in \underset{a' \in A}{\operatorname{argmax}} Ev\left(M_{\{a'\}}\right)\}$.
2. For $k = 2, \ldots, |A|$, let

$$a_{(k)} = \min\{a : a \in \underset{a' \in A \setminus \{a_{(1)}, \ldots, a_{(k-1)}\}}{\operatorname{argmax}} Ev\left(M_{\{a_{(1)}, \ldots, a_{(k-1)}, a'\}}\right)\},$$

where argmax is taken to be a set valued function. In the ordering of elements of A, ties are broken arbitrarily by taking the smaller element of $\{1, \ldots, p\}$. Thus, the ordered elements of A are tied to the labeling of covariates in $\mathbf{X}$ and any search algorithm that we define should be indifferent to relabeling.

Denote the ordered set of indices of $A = \{a_1, \ldots, a_{|A|}\}$ by $\mathring{A} = \left(a_{(1)}, \ldots, a_{(|A|)}\right)$. Define a partial ordering, $\trianglelefteq$, on $\mathcal{M}$ by setting $M_B \trianglelefteq M_A$ if and only if $B \subseteq A$ and $\mathring{A} = \left(\mathring{B}, a_{(|B|+1)}, \ldots, a_{(|A|)}\right)$ with equality if and only if $B = A$. We require that $M_0 \triangleleft M_A$ for all non-empty $A \subset \{1, \ldots, p\}$. To demonstrate this partial ordering, consider $p = 2$ and suppose that $Ev\left(M_{\{1\}}\right) \geq Ev\left(M_{\{2\}}\right)$. Then $M_0 \triangleleft M_{\{1\}} \triangleleft M_{\{1,2\}}$ and $M_0 \triangleleft M_{\{2\}}$, but $M_{\{2\}}$ is incomparable to $M_{\{1\}}$ or $M_{\{1,2\}}$.

This partial ordering, $\trianglelefteq$, can be extended to a total ordering, $\preceq$, by defining $M_B \preceq M_A$ if $M_B \trianglelefteq M_A$, $M_B = M_A$ if and only if $B = A$, and by specifying the following rule for comparing models that are incomparable under $\trianglelefteq$. Let $A = \{a_1, \ldots, a_{|A|}\}$ and $B = \{b_1, \ldots, b_{|B|}\}$. If M_A and M_B are incomparable under $\trianglelefteq$, then there exists a smallest integer $0 < k \leq \min\{|A|, |B|\}$ such that $a_{(k)} \neq b_{(k)}$. If $Ev\left(M_{\{a_{(1)}, \ldots, a_{(k)}\}}\right) > Ev\left(M_{\{b_{(1)}, \ldots, b_{(k)}\}}\right)$, then set $M_A \prec M_B$. If $Ev\left(M_{\{a_{(1)}, \ldots, a_{(k)}\}}\right) < Ev\left(M_{\{b_{(1)}, \ldots, b_{(k)}\}}\right)$, then set $M_B \prec M_A$. If $Ev\left(M_{\{a_{(1)}, \ldots, a_{(k)}\}}\right) = Ev\left(M_{\{b_{(1)}, \ldots, b_{(k)}\}}\right)$, then set $M_A \prec M_B$ if $a_{(k)} < b_{(k)}$ and $M_B \prec M_A$ if $b_{(k)} < a_{(k)}$. Once again, ties are broken arbitrarily based on the size of the index. To demonstrate this complete ordering, return again to the case of $p = 2$ with $Ev\left(M_{\{1\}}\right) \geq Ev\left(M_{\{2\}}\right)$. Then $M_0 \prec M_{\{1\}} \prec M_{\{1,2\}} \prec M_{\{2\}}$.

To better demonstrate $\trianglelefteq$ and $\preceq$, we consider a model space with $p = 4$ covariates; Figure 1 contains a possible partial ordering (left panel) and complete ordering (right panel) of the elements

of the model space. In both panels of Figure 1, there are eight rows and five columns with $\emptyset$ in position $(1, 1)$.

We describe the partial ordering $\trianglelefteq$ using three rules that describe the arrows in the left panel of Figure 1. Consider the model M in row i and column j. If $(i, j) \neq (1, 1)$ and $(i, j) \neq (2^{p-1}, 2)$, then there is a smallest row number $i' > i$ such that there is a model in cell (i', j) or $(i', j-1)$. Now we describe the rules for model M. First, M is $\triangleleft$ the model in cell (k, ℓ) for $i \leq k < i'$ and $j < \ell \leq p+1$. Second, M is incomparable to all models in row k for all $k \geq i'$. Third, M_0 (represented by $\emptyset$) is $\triangleleft$ all other models and the model in cell $(2^{p-1}, 2)$ (in this case, $M_{\{4\}}$) is not $\triangleleft$ any model. The model M in position (i, j) is $\prec$ the models in cells (i, ℓ) for $\ell > j$ and the models in cells (k, ℓ) for $k > i$ and $\ell > 1$.

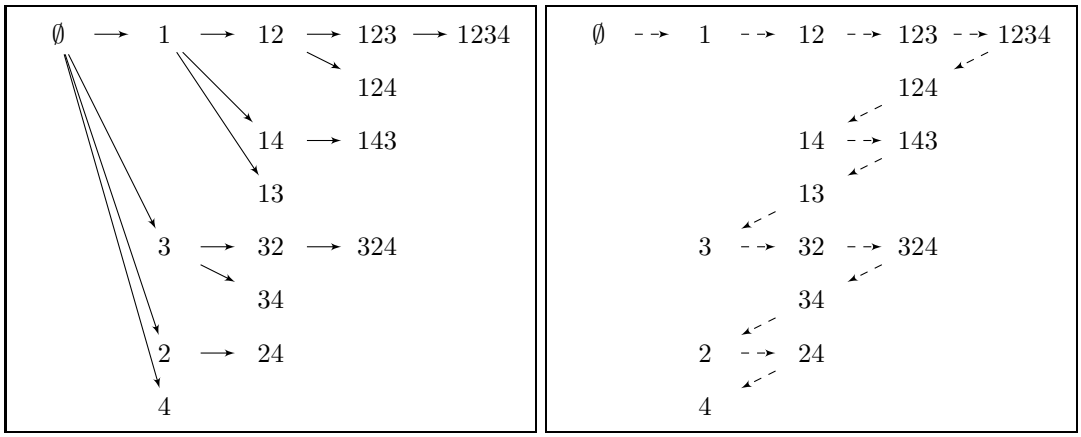

Fig. 1. Possible ordering of models on four covariates. Models are represented by ordered strings of covariate indices. The solid arrows on the left panel depict the partial increasing ordering $\triangleleft$ of the models and the dashed arrows on the right panel depict the increasing total ordering $\prec$ of the models. The symbol $\emptyset$ denotes the set of covariates in the model M_0.

The model space can be enumerated using $\prec$ by reading left to right and top to bottom in the right panel of Figure 1. A recursive algorithm for filling the model space can be described thusly. First, at $\emptyset$ the elements of $\{1, \ldots, p\}$ are ordered using $\prec$ for models with one covariate, which fills column 2. We label these elements $a_1 = 1, a_2 = 3, a_3 = 2, a_4 = 4$. This leads to p reduced model spaces corresponding to $\mathcal{A}_1 = \{A \in \mathcal{M} : a_1 \in A\}$ and $\mathcal{A}_i = \{A \in \mathcal{M} : a_i \in A,\ a_j \notin A \forall\, j = 1, \ldots, i-1\}$ for $i = 2, \ldots, p$. Each $\mathcal{A}_i$ corresponds to a smallest model $A_i = \{a_i\}$ and a set of covariates that can be added to it $S_i = \{a_{i+1}, \ldots, a_p\}$ for $i = 1, \ldots, p-1$ with $|S_i| = p - i$.

The filling algorithm moves to $\mathcal{A}_1$ and orders the elements in S_1 by appending them to A_1 and using $\prec$. This creates the first four rows in column 3 with models given by $\{1, 2\} \prec \{1, 4\} \prec \{1, 3\}$ with associated sets of potentially added indices $\{3, 4\}, \{3\}, \{\}$. The algorithm then keeps moving

to the right in row one, creating models and sets of potentially added indices. When the last model in a row is reached, the algorithm moves to the next row. For instance, once $\{3, 4\}$ is reached, the algorithm moves to the model $\{2\}$ with the set of potentially added indices $\{4\}$ and the algorithm continues moving across row seven. This process is completed until the model space is fully enumerated.

4.2. Pruning Rules

A pruning rule is any rule that allows us to remove large sets of models at once from the model space without considering each model individually. At any point in the filling algorithm described for Figure 1, the process can be interrupted and instead of moving to the right, the pruning forces the algorithm to move back one column and down that column to the next available model, which will come associated with a set of indices for potential inclusion. The pruned filling algorithm concludes when a pruning occurs in column 2 of Figure 1 or when the model in position $(2^{p-1}, 2)$ is reached.

The first pruning rule is a local rule wherein $0 < \epsilon_L < 1$ will control whether a covariate should be added to a model. For each model M_A, there is a unique model $P(M_A)$, referred to as the parent of M_A, such that $P(M_A)$ is the largest model M with $M \lhd M_A$. Define $Q_L(M_A) = Ev(M_A)/(Ev(M_A)+Ev(P(M_A)))$ to be the quality M_A normalized to the restricted set $\{M_A, P(M_A)\}$. The step to M_A is taken so long as $Q_L(M_A) > \epsilon_L$. One issue with the local rule is that it might be difficult to control individual false positives due to the large number of covariates. This will be especially difficult when proposing a model M_A that only contains covariates with true regression coefficient zero. A new "noisy" covariate is likely to be added to the parent model $P(M_A)$ due to fact that the learning rate of the Bayes factor is only a power of n. One advantage of the local rule is that it treats ties in the same manner, and is thus indifferent to relabeling of the covariates.

The second is a global rule, wherein M_A is not compared to its parent, but rather to the globally best model that has been visited thus far. Denote this model by M_G, though it could change at any step in the algorithm. Let $0 < \epsilon_G < 1$ and define $Q_G(M_A) = Ev(M_B)/(Ev(M_G) + Ev(M_A))$. The step to M_A is taken so long as $Q_G(\nu') > \epsilon_G$. Note that the M_G when proposing a step to model M_A is uniquely determined for each M_A once the decision rule is in place. The issue with the global rule is that it will be difficult or impossible to move to a model M_A that uses less covariates than M_G. Thus, this rule prunes the model space too aggressively and good models might be accidentally missed. Moreover, the global rule is that it is not indifferent to ties. Additionally, an alternative

global rule that uses the proposed model and all previously visited models is similarly ineffectual at moving to models with smaller dimension than M_G.

The third stopping rule is a path based stopping rule. Let $0 < \epsilon_P < 1$ and define $Q_P(M_A) = Ev(M_A)/\left(\sum_{M_B \trianglelefteq M_A} Ev(M_B)\right)$. The step to M_A is taken so long as $Q_P(M_A) > \epsilon_P$. Covariates are added until the proposed model has a sufficiently small posterior probability relative to its predecessors by $\lhd$. This allows addition of covariates while acknowledging that false positives may occur and that a penalization greater than just power of n is necessary to control these. Like the local rule, this pruning rule is indifferent to ties.

Now, we describe the pruning rules using Figure 1. The local rule considers comparing a model to its immediate predecessor. For example $\{1,2,3\}$ is compared to $\{1,2\}$ and $\{1,3\}$ is compared to $\{1\}$. The global rule compares a model to the global best that has already been visited. For example, suppose that $\{1,3,5\}$ is the posterior mode and that the best model visited before $\{1,3,5\}$ is $\{1,2\}$. If $\{1,2\}$ dramatically better than $\{1,3\}$ and the step to $\{1,3\}$ is rejected, then $\{1,3,5\}$ would never be visited. The path based rule compares a model to the models in the path before it. For example, $\{3,2,5\}$ is compared to the set of models $\{\emptyset,\{3\},\{3,2\}\}$ and $\{1,2,5,4\}$ is compared to the set of models $\{\emptyset,\{1\},\{1,2\},\{1,2,5\}\}$. Both the local rule and the path based rule can be defined in terms $\mathring{A}$, the set of the ordered set of elements from A.

5. Application of Methodology

This section compares the multiplicity correcting prior (Scott and Berger, 2010) and the complexity penalizing priors (Wilson et al., 2010; Castillo et al., 2015) to the limiting Poisson prior through the analysis of the diabetes dataset (Efron et al., 2004). The data contains 442 rows and 64 columns. An additional 372 variables are appended to the dataset. For columns 1 through 62 of the original data, six new columns are added. These columns have correlation which alternate sign over the 62 columns and whose absolute value is $i/\sqrt{1+i^2}$ for $i = 1,\ldots,6$. This results in $p = 436$ covariates, many of which are highly correlated with the original covariates. There are $2^{436} \approx 1.77 \times 10^{131}$ models. The priors considered are the descendant based prior with $\eta = 1$ and the Beta-Binomial(a,b) priors for $a = 1$ and $b \in \{1, p, p^{1.5}\}$. The prior for each model is taken to be the intrinsic prior and the algorithm used is the path-based pruning rule with $\epsilon_P = 0.01$.

The models are compared in terms of both their posterior concentration and the amount of uncertainty that they reflect in the model space. This is the delicate trade off between penalizing enough to get a meaningful posterior and penalizing too much and receiving a posterior that concentrates too quickly. Table 1 provides a summary of information for the posterior on the

Table 1. Posterior summary for the analysis of the diabetes data with added covariates. D($\eta = 1$) is the descendant prior with $\eta = 1$ and B-B(a, b) is the Beta-Binomial prior with parameters (a, b).

Prior	D($\eta = 1$)	B-B(1, 1)	B-B(1, p)	B-B(1, $p^{1.5}$)
Mode	$3, 9$	$3, 4, 9$	$3, 4, 9$	$3, 9$
P(Mode)	0.405	0.065	0.160	0.820
# of Models Considered	15,423	32,018	21,427	8,192
# of Models s/t P(Models)> 0.9	51	5,409	1,579	2
# of covariates with $P(\gamma = 1) > 0.1$	4	8	5	3
$P(\gamma_3 = 1)$	1.000	1.000	1.000	1.000
$P(\gamma_4 = 1)$	0.381	0.852	0.666	0.105
$P(\gamma_9 = 1)$	1.000	1.000	1.000	1.000
$P(\gamma_{338} = 1)$	0.126	0.483	0.296	0.028

model space. Across the columns are the four considered priors for the model space: descendant with $\eta = 1$, denoted by D($\eta = 1$), and then the Beta-Binomial(a, b), denoted by B-B(a, b). In the rows, the table catalogs the posterior probability of the modal model, the variables in the modal model, the number of models considered during the search algorithm, the number of models necessary to comprise 90% of the models proposed, the number of covariates with posterior inclusion probability greater than 10%, and information on interesting covariates. One particular instance is covariate 338, which was artificially created and should be considered a false covariate, but has non-negligible posterior inclusion probability for all but the most penalizing prior.

As can be seen in Table 1, the descendant prior, which produces a limiting Poisson prior on the model space, produces the most satisfactory results in this analysis. First, consider posterior probability of the highest probability model. The descendant prior places 40% probability on model $\{3, 9\}$, which provides strong, but not conclusive, evidence that this is the true model. The Beta-Binomial prior with $b = p^{1.5}$ reflects its stronger penalization of covariate inclusion, placing over 80% of the posterior probability on $\{3, 9\}$, overstating the evidence in favor of this model. In contrast, the Beta-Binomial priors with $b \in \{1, p\}$ produce a much smaller amount of mass on $\{3, 4, 9\}$ as the highest posterior probability model, reflecting their poor penalization of additional covariates. This is most drastic when $b = 1$ and under 7% of the posterior is placed on $\{3, 4, 9\}$. The penalization of the priors is also reflected in the number of models that it takes to comprise 90% of the posterior probability of the models considered. The descendant prior requires 52 models. This reflects reasonable uncertainty in the posterior of the model space, while concentrating well

on good models. In contrast, the Beta-Binomial with $b = p^{1.5}$ only requires two models for over 90% posterior probability. The concentration of the posterior is too strong and the certainty in the top few models is overstated. The Beta-Binomial priors with $b \in \{1, p\}$ do not concentrate fast enough and require an unreasonable number of models to achieve 90% posterior probability. This is most pronounced for $b = 1$, which requires over $5,000$ models.

Finally, while all of the model space priors provide certainty that covariates 3 and 9 should be included in the model, the inclusion probabilities for covariates 4 and 338 are dramatically different. The descendant prior produces weak evidence against the inclusion of covariate 4 and strong evidence against the inclusion of covariate 338. The Bayes factor of $M_{\{3,4,9\}}$ vs $M_{\{3,9\}}$ is 351.84, which seems to provide sufficient evidence for the inclusion of 4. However, this same logic could be applied to the model $M_{\{3,9,338\}}$ whose Bayes factor against $M_{\{3,9\}}$ is 86.83. In order to keep the false covariate 338 from entering the model, the descendant prior also penalizes that most likely true covariate 4, but not excessively. In contrast, the Beta-Binomial with $b = p^{1.5}$ penalizes too much and places too small of a probability on the inclusion of covariate 4. Certainly, the Beta-Binomial with $b \in \{1, p\}$ produce posterior inclusion probabilities for covariate 338. Only the descendant prior produces a posterior that provides a good balance of Type I and Type II error rates in nested comparisons.

For completeness, a comparison is made to non-local priors using the default settings of the `R` package `mombf`. The posterior inclusion probabilities for covariates $3, 4, 9, 338$ are $1, 0.284, 1, 0.071$, respectively. These probabilities compare well to those obtained by the descendant prior on the model space with intrinsic prior on the model specific parameters, though the former are smaller than the latter for covariates 4 and 338.

6. Discussion

This paper has presented both theory and methods for performing Bayesian regression in sparse settings. The key is to provide a prior on the model space with both a meaningful interpretation and a construction from first principles. The priors presented here rely on a reformulation of the way in which the model space is viewed. Instead of considering indicator variables for covariate inclusion, the model space is viewed as a series of local null and alternative hypotheses. In the limit when the number of covariates increases to ∞, a self similarity argument suggests that the Poisson is the only reasonable limiting prior on the number of covariates in the model. Two constructions were given that provide such a limiting prior. Using the proposed model space prior, consistency for a finite true model is shown under the same conditions as Johnson and Rossell (2012), but with

the prior on model specific parameters being the intrinsic prior.

The expansion of the paper to the case where the mean structure is given by and infinite series or when $p >> n$ are active areas of research. There are a a few difficulties in producing reasonable objective Bayesian inferences in these frameworks. First, the choice of the prior on the model specific parameters has a pronounced effect when the size of the model is large. Changing $\pi(w)$ to a Gamma$(0.5, 0.5)$ produces the Zellner-Siow prior, which concentrates more at the origin than does the intrinsic prior and helps to identify weak signals. Similarly, choosing a more flexible family can provide even more shrinkage (Liang et al., 2008; Womack et al., 2014). Consistency of pairwise Bayes factors or posterior odds is delicate when using these priors and simple approximations are mostly useless when comparing a true model with an infinite number of covariates to a model with a finite number more covariates. The asymptotic behavior of the Bayes factors has to be properly balanced with the penalization provided by the model space prior, as discussed in Castillo et al. (2012) and Castillo et al. (2015).

Second, the objective Bayes framework using scaled mixtures of g-priors becomes problematic when $p > n$. Similarly, the conditions on the design matrix become more delicate. For the objective Bayes framework, it seems natural to discuss models not only in terms of covariate inclusion, but also in terms of the subspace of the outcome that is being modeled. The assumptions in this paper forced these to be in a 1-1 correspondence without too much correlation between modeled subspaces. When $p > n$, it seems natural to consider a model as being defined solely by the subspace of the outcome which is being modeled. A natural sieve is created because any model with $k \geq n$ covariates that is rank n interpolates the data. The implications of this for both the Bayes factors and the model space prior need to be dealt with carefully.

Third, one should be able to extend all of the results of Castillo et al. (2015) to the objective Bayes framework, replacing their Assumption 1 (see (5)) with the limiting Poisson prior and using an appropriate scaled mixture of g-priors . The limiting Poisson prior on the number of covariates provides an adaptive version of their Assumption 1, which it satisfies for large enough s — see Example 2.3 in Castillo et al. (2012). One difficulty for implementation of the methods in Castillo et al. (2015) is the use of independent Laplace priors or heavy-tailed versions thereof, which make computation of the Bayes factor difficult when the dimension of the model is large. The advantage of scaled mixtures of g-priors is that the Bayes factor can be computed as a simple one dimensional integral that only depends on the dimension of the model and the statistic R^2. Simple bijective transformations of w allow the integral in (12) to be computed within machine error for large model dimension and R^2 close to one.

Finally, the algorithm outlined in this paper allows us to take advantage of the developedprior on the model space. The path-based algorithm, in conjunction with the model space prior, provide for easy learning of good models while not foraging around too many bad models. In contrast to MCMC or RJMCMC methods, the proposed algorithm will never get stuck in a local mode and never revisits the same model twice. Similarly, the renormalized posterior on the set of considered models is a much better approximation to the posterior than the site-frequency analysis provides. The gains in efficiency are staggering, exploring the posterior for the diabetes example with the descendant prior, path-based rule, and $\epsilon_P = 0.01$ takes under three seconds.

References

Altomare, D., Consonni, G., and La Rocca, L. (2013). Objective bayesian search of gaussian directed acyclic graphical models for ordered variables with non-local priors. *Biometrics*, 69(2):478–487.

Berger, J. O. and Pericchi, L. R. (1996). The intrinsic bayes factor for model selection and prediction. *Journal of the American Statistical Association*, 91(433):109–122.

Casella, G., Girón, F. J., Martínez, M. L., and Moreno, E. (2009). Consistency of bayesian procedures for variable selection. *Annals of Statistics*, pages 1207–1228.

Casella, G. and Moreno, E. (2012). Objective bayesian variable selection. *Journal of the American Statistical Association*.

Castillo, I., Schmidt-Hieber, J., and van der Vaart, A. (2015). Bayesian linear regression with sparse priors. *Annals of Statistics*, 43(5):1986–2018.

Castillo, I., van der Vaart, A., et al. (2012). Needles and straw in a haystack: Posterior concentration for possibly sparse sequences. *Annals of Statistics*, 40(4):2069–2101.

Efron, B., Hastie, T., Johnstone, I., Tibshirani, R., et al. (2004). Least angle regression. *Annals of Statistics*, 32(2):407–499.

Girón, F. J., Moreno, E., Casella, G., and Martínez, M. (2010). Consistency of objective bayes factors for nonnested linear models and increasing model dimension. *RACSAM-Revista de la Real Academia de Ciencias Exactas, Fisicas y Naturales. Serie A. Matematicas*, 104(1):57–67.

Hans, C. (2010). Model uncertainty and variable selection in bayesian lasso regression. *Statistics and Computing*, 20(2):221–229.

Johnson, V. E. and Rossell, D. (2010). On the use of non-local prior densities in bayesian hypothesis tests. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(2):143–170.

Johnson, V. E. and Rossell, D. (2012). Bayesian model selection in high-dimensional settings. *Journal of the American Statistical Association*, 107(498):649–660.

Knight, K. and Fu, W. (2000). Asymptotics for lasso-type estimators. *Annals of Statistics*, pages 1356–1378.

Ley, E. and Steel, M. F. (2009). On the effect of prior assumptions in bayesian model averaging with applications to growth regression. *Journal of Applied Econometrics*, 24(4):651–674.

Liang, F., Paulo, R., Molina, G., Clyde, M. A., and Berger, J. O. (2008). Mixtures of g priors for bayesian variable selection. *Journal of the American Statistical Association*, 103(481).

Moreno, E., Girón, F. J., and Casella, G. (2010). Consistency of objective bayes factors as the model dimension grows. *Annals of Statistics*, pages 1937–1952.

Scott, J. G. and Berger, J. O. (2010). Bayes and empirical-bayes multiplicity adjustment in the variable-selection problem. *Annals of Statistics*, 38(5):2587–2619.

Wilson, M. A., Iversen, E. S., Clyde, M. A., Schmidler, S. C., and Schildkraut, J. M. (2010). Bayesian model search and multilevel inference for snp association studies. *Annals of Applied Statistics*, 4(3):1342.

Womack, A. J., León-Novelo, L., and Casella, G. (2014). Inference from intrinsic bayes' procedures under model selection and uncertainty. *Journal of the American Statistical Association*, 109(507):1040–1053.

Zellner, A. (1986). On assessing prior distributions and bayesian regression analysis with g-prior distributions. *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno De Finetti*, 6:233–243.

Zellner, A. and Siow, A. (1980). Posterior odds ratios for selected regression hypotheses. *Trabajos de Estadística y de Investigación Operativa*, 31(1):585–603.